

**Exam : JN0-664**

**Title : Service Provider  
Professional (JNCIP-SP)**

**<https://www.passcert.com/JN0-664.html>**

1. Which two statements are correct about reflecting inet-vpn unicast prefixes in BGP route reflection? (Choose two.)


- A. Route reflectors do not change any existing BGP attributes by default when advertising routes.
- B. A BGP peer does not require any configuration changes to become a route reflector client.
- C. Clients add their originator ID when advertising routes to their route reflector
- D. Route reflectors add their cluster ID to the AS path when readvertising client routes.

**Answer:** A,B

**Explanation:**

Route reflection is a BGP feature that allows a router to reflect routes learned from one IBGP peer to another IBGP peer, without requiring a full-mesh IBGP topology. Route reflectors do not change any existing BGP attributes by default when advertising routes, unless explicitly configured to do so. A BGP peer does not require any configuration changes to become a route reflector client, only the route reflector needs to be configured with the client parameter under [edit protocols bgp group group-name neighbor neighbor-address] hierarchy level.

2. Exhibit



Click the Exhibit button-Referring to the exhibit, which two statements are correct about BGP routes on R3 that are learned from the ISP-A neighbor? (Choose two.)

- A. By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3.
- B. The BGP local-preference value that is used by ISP-A is not advertised to R3.
- C. All BGP attribute values must be removed before receiving the routes.
- D. The next-hop value for these routes is changed by ISP-A before being sent to R3.

**Answer:** A,B

**Explanation:**

BGP is an exterior gateway protocol that uses path vector routing to exchange routing information among autonomous systems. BGP uses various attributes to select the best path to each destination and to propagate routing policies. Some of the common BGP attributes are AS path, next hop, local preference, MED, origin, weight, and community. BGP attributes can be classified into four categories: well-known mandatory, well-known discretionary, optional transitive, and optional nontransitive. Well-known mandatory attributes are attributes that must be present in every BGP update message and must be recognized by every BGP speaker. Well-known discretionary attributes are attributes that may or may not be present in a BGP update message but must be recognized by every BGP speaker. Optional transitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional transitive attribute is not recognized by a BGP speaker, it is passed along to the next BGP speaker. Optional nontransitive attributes are attributes that may or may not be present in a BGP update message and may or may not be recognized by a BGP speaker. If an optional nontransitive attribute is not recognized by a BGP speaker, it is not passed along to the next BGP speaker. In this question, we have four routers (R1, R2, R3, and R4) that are connected in a full mesh topology and running IBGP. R3 receives the 192.168.0.0/16 route from its EBGP neighbor and advertises it to R1 and R4 with different BGP attribute values. We are asked which statements are correct about the BGP routes on R3 that are learned from the ISP-A neighbor.

Based on the information given, we can infer that the correct statements are:

- ☞ By default, the next-hop value for these routes is not changed by ISP-A before being sent to R3. This is because the default behavior of EBGP is to preserve the next-hop attribute of the routes received from another EBGP neighbor. The next-hop attribute indicates the IP address of the router that should be used as the next hop to reach the destination network.
- ☞ The BGP local-preference value that is used by ISP-A is not advertised to R3. This is because the local-preference attribute is a well-known discretionary attribute that is used to influence the outbound traffic from an autonomous system. The local-preference attribute is only propagated within an autonomous system and is not advertised to external neighbors.

**References:**

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html>:  
<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13762-40.html>:  
<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html>

3.A packet is received on an interface configured with transmission scheduling. One of the configured queues.

In this scenario, which two actions will be taken by default on a Junos device? (Choose two.)

- A. The excess traffic will be discarded
- B. The exceeding queue will be considered to have negative bandwidth credit.
- C. The excess traffic will use bandwidth available from other queues
- D. The exceeding queue will be considered to have positive bandwidth credit

**Answer:** A,B

**Explanation:**

Transmission scheduling is a CoS feature that allows you to allocate bandwidth among different queues on an interface. Each queue has a configured bandwidth percentage that determines how much of the available bandwidth it can use. If a queue exceeds its allocated bandwidth, it is considered to have

negative bandwidth credit and its excess traffic will be discarded by default. If a queue does not use all of its allocated bandwidth, it is considered to have positive bandwidth credit and its unused bandwidth can be shared by other queues.

4. Which two statements are correct about VPLS tunnels? (Choose two.)


- A. LDP-signaled VPLS tunnels only support control bit 0.
- B. LDP-signaled VPLS tunnels use auto-discovery to provision sites
- C. BGP-signaled VPLS tunnels can use either RSVP or LDP between the PE routers.
- D. BGP-signaled VPLS tunnels require manual provisioning of sites.

**Answer:** B,C

**Explanation:**

VPLS is a Layer 2 VPN technology that allows multiple sites to connect over a shared IP/MPLS network as if they were on the same LAN. VPLS tunnels can be signaled using either Label Distribution Protocol (LDP) or Border Gateway Protocol (BGP). LDP-signaled VPLS tunnels use auto-discovery to provision sites, meaning that PE routers can automatically discover other PE routers that belong to the same VPLS instance

5. Exhibit



The network shown in the exhibit is based on IS-IS

Which statement is correct in this scenario?

- A. The NSEL byte for Area 0001 is 00.
- B. The area address is two bytes.
- C. The routers are using unnumbered interfaces
- D. The system ID of R1\_2 is 192.168.16.1



**Answer: A**

**Explanation:**

IS-IS is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. IS-IS uses two types of addresses to identify routers and areas: system ID and area address. The system ID is a unique identifier for each router in an IS-IS domain. The system ID is 6 octets long and can be derived from the MAC address or manually configured. The area address is a variable-length identifier for each area in an IS-IS domain. The area address can be 1 to 13 octets long and is composed of high-order octets of the address. An IS-IS instance may be assigned multiple area addresses, which are considered synonymous. Multiple synonymous area addresses are useful when merging or splitting areas in the domain<sup>1</sup>. In this question, we have a network based on IS-IS with four routers (R1\_1, R1\_2, R2\_1, and R2\_2) belonging to area 0001. The area address for area 0001 is 49.0001. The NSEL byte for area 0001 is the last octet of the address, which is 01. The NSEL byte stands for Network Service Access Point Selector (NSAP Selector) and indicates the type of service requested from the network layer<sup>2</sup>. Therefore, the correct statement in this scenario is that the NSEL byte for area 0001 is 01.

References:


1:

[https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute\\_isis/configuration/xr-16/irs-xe-16-book/irs-ovrv-w-cf.html](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_isis/configuration/xr-16/irs-xe-16-book/irs-ovrv-w-cf.html)

2:

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html>

**6.Exhibit**



Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts

Which two statements are correct in this scenario? (Choose two.)

- A. You must configure a static route in the main routing instance for the 10.1.2.0/24 prefix that uses the VPN-A.inet.0 table as the next hop
- B. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-A.inet.0 table as the next hop.

C. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN-A.inet.0 table.

D. You must configure a RIB group on PE-1 to leak the 10.1.2.0/24 prefix from the VPN-A.inet.0 table to the inet.0 table.

**Answer:** A,B

**Explanation:**

To provide Internet access for the hosts at Site 1, you need to configure static routes in the main routing instance on PE-1 that point to the VPN-A.inet.0 table as the next hop. This allows PE-1 to forward traffic from the Internet to CE-1 using MPLS labels and vice versa. You need to configure two static routes: one for the 10.1.2.0/24 prefix that represents the private network of Site 1, and one for the 203.0.113.1/32 prefix that represents the public IP address of CE-1.

7. Which three mechanisms are used by Junos platforms to evaluate incoming traffic for CoS purposes? (Choose three)

- A. rewrite rules
- B. behavior aggregate classifiers
- C. traffic shapers
- D. fixed classifiers
- E. multifield classifiers

**Answer:** B,D,E


**Explanation:**

Junos platforms use different mechanisms to evaluate incoming traffic for CoS purposes, such as:

- ☞ Behavior aggregate classifiers: These classifiers use a single field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined or user-defined values.
- ☞ Fixed classifiers: These classifiers use a fixed field in a packet header to classify traffic into different forwarding classes and loss priorities based on predefined values.
- ☞ Multifield classifiers: These classifiers use multiple fields in a packet header to classify traffic into different forwarding classes and loss priorities based on user-defined values and filters.

Rewrite rules and traffic shapers are not used to evaluate incoming traffic for CoS purposes, but rather to modify or shape outgoing traffic based on CoS policies.

8. Exhibit



You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number.

Which two steps will accomplish this task? (Choose two.)

- A. Configure the BGP group with the advertise-peer-as parameter on R1 and R4.
- B. Configure the BGP group with the as-override parameter on R2 and R3
- C. Configure the BGP group with the advertise-peer-as parameter on R2 and R3.
- D. Configure the BGP group with the as-override parameter on R1 and R4

**Answer:** A,B

**Explanation:**

The advertise-peer-as parameter allows a router to advertise its peer's AS number as part of the AS path attribute when sending BGP updates to other peers. This parameter is useful when two routers in the same AS need to exchange routes through another AS, such as in the case of R1 and R4. By configuring this parameter on R1 and R4, they can advertise each other's AS number to R2 and R3, respectively.

The as-override parameter allows a router to replace the AS number of its peer with its own AS number when receiving BGP updates from that peer. This parameter is useful when two routers in different ASes need to exchange routes through another AS that has the same AS number as one of them, such as in the case of R2 and R3. By configuring this parameter on R2 and R3, they can override the AS number of R1 and R4 with their own AS number when sending BGP updates to each other.

9.You want to ensure that L1 IS-IS routers have only the most specific routes available from L2 IS-IS routers.

Which action accomplishes this task?

- A. Configure the ignore-attached-bit parameter on all L2 routers.
- B. Configure all routers to allow wide metrics.
- C. Configure all routers to be L1.
- D. Configure the ignore-attached-bit parameter on all L1 routers

**Answer: D**

**Explanation:**

The attached bit is a flag in an IS-IS LSP that indicates whether a router is connected to another area or level (L2) of the network. By default, L2 routers set this bit when they advertise their LSPs to L1 routers, and L1 routers use this bit to select a default route to reach other areas or levels through L2 routers.

However, this may result in suboptimal routing if there are multiple L2 routers with different paths to other areas or levels. To ensure that L1 routers have only the most specific routes available from L2 routers, you can configure the ignore-attached-bit parameter on all L1 routers. This makes L1 routers ignore the